

2025-09-19

The self-archived post print version of this conference paper is available at Linköping University Institutional Repository (DiVA):

<https://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-217838>

Exploring the Properties of Multi-Agent Terrain-Aided Navigation

Eric Sevonius, Fredrik Gustafsson and Gustaf Hendeby

In: 2025 28th International Conference on Information Fusion (FUSION), 2025, pp. 1-8.

ISBN: 9781037056239

Publisher: IEEE

<https://doi.org/10.23919/FUSION65864.2025.11124164>

N.B.: When citing this work, cite the original publication.

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Exploring the Properties of Multi-Agent Terrain-Aided Navigation

Eric Sevonius^{*†}, Fredrik Gustafsson[†], and Gustaf Hendeby[†]

^{*} Saab Dynamics, Linköping, Sweden

e-mail: eric.sevonius1@saabgroup.com

[†] Dept. of Electrical Engineering, Linköping University, Linköping, Sweden

e-mail: firstname.lastname@liu.se

Abstract—Due to recent events that have demonstrated the vulnerabilities of *global navigation satellite systems* (GNSS) there has been an increased interest in alternative methods for localization. One traditional alternative method is *terrain-aided navigation* (TAN), where a platform localizes itself by measuring the terrain elevation and comparing it to a *digital elevation map* (DEM). While single-agent TAN has been extensively studied, multi-agent TAN remains less explored. This paper addresses the multi-agent TAN problem with a focus on its properties. We formulate a *weighted least squares* (WLS) estimator for computing a snapshot solution to the problem and formulate a *Cramér-Rao Lower Bound* (CRLB) to evaluate it. Using the expressions for the estimator and the CRLB we are able to highlight some insightful properties of the problem. The findings are verified in a simulation study where we evaluate the performance with respect to the altitude sensor accuracy, the group formation accuracy, the number of agents and their formation. Notably, we observe that the solution is relatively insensitive to errors in agent position, suggesting that low-accuracy inertial navigation systems and distance sensors are sufficient for determining their positions. Increasing the number of agents beyond a few seems to have a large effect on both the efficiency and robustness of the estimator, which lessens as the number of agents increases. However, increasing the number of agents does not compensate for poor altitude sensor quality. Additionally, while spatial separation between agents is important for effective map utilization, further separation beyond a certain point does not enhance performance. These findings provide design guidelines for multi-agent TAN systems and identify areas for further research.

Index Terms—Sensor Fusion, Positioning, Terrain-Aided Navigation, Multi-Agent

I. INTRODUCTION

Recent events have shown that *global navigation satellite systems* (GNSS) are susceptible to interference by malicious actors [1]. However, the need for high precision localization remains. One way to meet this need is to look for alternative methods utilizing other sources of information. One class of information sources predating GNSS are natural signals. Natural signals are features in the environment that can be mapped and used for localization, examples include terrain elevation [2], magnetic anomalies [3], and the patterns of the stars [4].

This work was performed within the Competence Center SEDDIT (Sensor Informatics and Decision making for the DIgital Transformation), supported by Sweden's Innovation Agency within the research and innovation program Advanced digitalization.

Navigation using terrain elevation, commonly referred to as *terrain-aided navigation* (TAN) [5–7], has been explored extensively. Applications include airplanes measuring the terrain elevation using radar [2], *autonomous underwater vehicles* (AUVs) using sonar to measure the bottom depth [8], and even ground vehicles using inertial sensors to estimate their vertical displacement, and in turn the terrain height [9]. An inherent difficulty in utilizing measurements of the elevation of the terrain is that a single measurement will generally correspond to multiple points in an elevation database, even over a small area. By collecting multiple measurements of the elevation at approximately known locations, the number of possible locations can ultimately be narrowed down significantly. This is typically achieved by moving across the terrain and estimating the relative displacement between the measurements using inertial sensors [7].

Recently, cooperative localization has seen increased interest [10–12]. A group of platforms equipped with sensors can leverage the information from its neighbors to improve the overall positioning of the whole group. Instead of only sampling the map in time, multiple agents can act as a sensor array and sample spatially in TAN. This is conceptually similar to AUVs employing multibeam sonar [8], but with a larger footprint. Some work on the topic of multi-agent TAN has been done [13–15]. In [14] a framework for cooperative localization using TAN is presented. The results from a simulated experiment using real terrain maps show that cooperative TAN outperforms traditional TAN both in terms of positional accuracy and robustness.

In this paper, instead of focusing on the practical aspects of fusing the information in a decentralized fashion, we examine what properties affect performance. Specifically, we are interested in determining how the quality of altimeter sensors and relative distance estimates affect the performance of the estimator. These are factors that can be affected through, *e.g.*, using more expensive sensors. Additionally, we examine how the composition of the array affects performance by considering the effect of varying the number of sensors as well as the mean distance between them.

The work is organized as follows. Sec. II describes the problem description considered in the paper. Sec. III presents the proposed estimator. In Sec. IV the *Cramér-Rao Lower Bound* (CRLB) for the problem is specified. Sec. V discusses

the representation of the map. Sec. VI presents the simulation setup and simulation results. Finally, in Sec. VII the conclusions from the results are discussed.

II. PROBLEM DESCRIPTION

This paper considers the problem of TAN, where measurements of the environment are used to estimate the position of the sensor. The general problem is to estimate the trajectory taken by the sensor, as it moves over time, which leads to a state-space estimation problem. This section establishes the general estimation problem formulation this paper addresses, and then extends the often used single sensor terrain measurement model to more general multi-sensor models, which will be analyzed in this paper.

A. Navigation Model

The goal is to use terrain measurements in a navigation filter, defined by the model

$$x_{k+1} = f(x_k, u_k, v_k), \quad (1a)$$

$$y_n = h(x_k) + e_k, \quad (1b)$$

where x_k represents the navigation state, u_k is known input such as odometry, y_k represents the available measurements, and v_k and e_k are the process and measurement noise, respectively, all at time k . The navigation state typically comprises the position, kinematics such as velocity and sometimes acceleration, and potential additional properties of interest. In this paper, the focus is terrain measurements; however additional measurements, *e.g.*, positions from a GPS system are possible.

For the remainder of this paper, the focus will be understanding the properties of the observations of the terrain. Without loss of generality, and to simplify notation, from here on x_k will be assumed to be the position. Furthermore, as we study measurements from a single time instance, the time indicator k will be dropped, too.

B. Terrain Measurement

1) *Single Sensor Terrain Measurement*: Considering a single sensor, with the simplified notation, the measurement equation becomes

$$y = h(x) + e, \quad (2)$$

where y is the measured altitude, and $h(x)$ is the height of the terrain at position x . The function $h(x)$ can be represented by a *digital elevation model* (DEM), in which case the measurement noise e represents both errors in the altitude measured by the sensor and errors in the DEM. In many cases the terrain h could be heavily nonlinear, requiring the usage of nonlinear methods to solve the resulting estimation problem.

The model (2), can be realized by flying platform at a known mean sea level (using, *e.g.*, a barometer combined with inertial sensors) measuring the terrain elevation with a wide lobe radar, as depicted in Fig. 1. Often, the measurement noise can be approximated as Gaussian, $e \sim \mathcal{N}(0, \sigma_e^2)$. However, as reported in [7], the terrain can interfere with the measurement, *e.g.*, a part of the radar is reflected on the tree tops and not




Fig. 1: Illustration of a TAN scenario. Here, the altimeter measures the distance from the platform to the terrain, \bar{y} , and the elevation map evaluated at the true position corresponds to the terrain height, $h(x)$. The measurement can be expressed as $\bar{y} = z - h(x)$, where z is the altitude of the platform with respect to the *mean sea level* (MSL). By using a barometer and INS to measure z independently, a simpler measurement model can be derived which excludes it, $y = h(x)$.

the ground. In this case, using a Gaussian mixture to represent the measurement error is often more appropriate [7].

It is also possible to use more advanced sensor configurations for single sensor TAN, and in this way simplify the localization problem. An example of this is the multi-beam sonar sensor used for underwater localization, studied in [8]. In this case y , becomes a vector of several distance measurements, which would then also be extracted from the terrain model h .

2) *Multiple Sensors Terrain Measurement*: The natural extension to this is to introduce more than one sensor, separated in space. This dynamic sensor array provides information from several positions in a neighborhood, which can then be used for localization.

Here, an array of N sensors located at the relative positions d_n , $n = 0, \dots, N - 1$, will be considered. The first sensor is considered to constitute the origin of the array, *i.e.*, $d_0 = 0$. This results in the individual measurements

$$y_n = h(x + d_n + w_n) + e_n, \quad (3)$$

where $w_n \sim \mathcal{N}(0, \Sigma_w)$ represents the uncertainty in the relative position in the array. Note that $n = 0$ should be treated slightly different from the other sensors, $w_n \equiv 0$ as a result of defining sensor $n = 0$ as the origin of the array. Estimating x will hence be the same as estimating the position of the first sensor, from which the other sensor locations can be derived.

The measurements in the array can now be stacked, $Y = [y_0 \ y_1 \ \dots \ y_{N-1}]^T$, and treated as a big measurement. Using the same stacked notation to obtain D , W , and E , the single sensor model structure is essentially preserved

$$Y = H(x \oplus (D + W)) + E$$

with

$$H(x \oplus (D + W)) = \begin{bmatrix} h(x) \\ h(x + d_1 + w_1) \\ \vdots \\ h(x + d_{N-1} + w_{N-1}) \end{bmatrix},$$

and where $x \oplus A$ is defined as adding the vector x to each vector a_n in the stacked vector A . Here, $E \sim \mathcal{N}(0, \Sigma_E)$ denotes the measurement errors, whose components are assumed mutually independent and identically distributed with $e_i \sim \mathcal{N}(0, \sigma_e^2)$ such that $\Sigma_E = \sigma_e^2 \cdot I$, and $W \sim \mathcal{N}(0, \Sigma_W)$ represents the uncertainty in the relative sensor positions D .

Throughout this paper, three different cases will be considered: (i) the array configuration is fully known, (ii) the uncertainty in all relative sensor positions are independent, and (iii) the relative sensor positions are uncertain with no further structure. The three cases represent increasingly difficult estimation problems.

The first case is what can be expected from sensors mounted on a rigid body. No additional uncertainty is introduced, and $W \equiv 0$, which can be interpreted as $\Sigma_W = 0$. This is essentially a special case of the single sensor case described in Sec. II-B1.

The second case requires that the relative positions of the sensors can be estimated independently, *e.g.*, using separate measurements to obtain all the relative position with respect to the first sensors. Assuming the uncertainty in the individual positions, $w_i \sim \mathcal{N}(0, \Sigma_{w,i})$ for $i \neq 0$, Σ_W becomes a (block-) diagonal matrix with $\Sigma_{w,i}$ on the diagonal.

Case three is the most general problem formulation, making no assumption about the uncertainty of the relative sensor positions, *i.e.*, $\Sigma_W \succeq 0$. This is typically the case when a separate estimation algorithm is used to estimate the group formation and the relative distances d_n , and this allows for almost any way of obtaining the relative sensor positions.

III. ESTIMATION

Our goal is to find an estimate \hat{x} of the first sensor in the array. To start with, we consider the estimation problem in the context of a one-dimensional terrain instead of an m -dimensional one. Doing so leads to more intuitive expressions, and is a trick we will employ throughout the paper. The conclusions drawn using these simpler expressions can then be applied to the m -dimensional case without loss of generality.

A. Known Geometry

We start by formulating an estimator for the case of a known sensor array geometry (i).

1) *Correlation Approach*: In the context of a one-dimensional terrain and sensors at known distances, we can consider the measurements by the formation as sampling a continuous time signal $h(x)$ at N discrete points. The standard approach in signal processing to find a signal hidden in noise is by matched filtering. The matched filter is a mirrored version

of the signal of interest. This coincides here with maximizing correlation. The estimate is defined by

$$\hat{x} = \arg \max_x \hat{R}_{yh}(x), \quad (4a)$$

where

$$\hat{R}_{yh}(x) = \frac{1}{N} \sum_{n=0}^{N-1} (y_n - \mu_y)(h(x + d_n) - \mu_h), \quad (4b)$$

$$\mu_y = \frac{1}{N} \sum_{n=0}^{N-1} y_n, \quad (4c)$$

$$\mu_h = \frac{1}{N} \sum_{n=0}^{N-1} h(x + d_n). \quad (4d)$$

If $h(x)$ is indeed a stationary stochastic process with covariance function $R_{hh}(s)$, then we have $\mathbb{E}(\hat{R}_{yh}(s)) = R_{hh}(s-x)$, which indicates that the faster decaying covariance function (“whiter map”), the better potential precision. However, the configuration $\{d_n\}_{n=0}^{N-1}$ does not show up in the first order moment $\mathbb{E}(\hat{R}_{yh}(s))$, but should show up in the variance $\text{var}(\hat{R}_{yh}(s))$.

2) *Least Squares Approach*: An alternative to the matched filter/correlation approach is to apply the *least squares* (LS) method. Since $Y \sim \mathcal{N}(H(x \oplus D), \Sigma_E)$ the LS estimate is given by

$$\hat{x} = \arg \min_x V(x), \quad (5a)$$

where

$$V(x) = \sum_{n=0}^{N-1} (y_n - h(x + d_n))^2. \quad (5b)$$

If we expand the LS loss function (5b), we get

$$\begin{aligned} V(x) &= \sum_{n=0}^{N-1} ((y_n - \mu_h) - (h(x + d_n) - \mu_h))^2, \\ &= \sum_{n=0}^{N-1} ((y_n - \mu_h)^2 + (h(x + d_n) - \mu_h)^2 \\ &\quad - 2(y_n - \mu_h)(h(x + d_n) - \mu_h)) \\ &= -2N\hat{R}_{yh}(x) + \sum_{n=0}^{N-1} ((y_n - \mu_h)^2 + (h(x + d_n) - \mu_h)^2). \end{aligned}$$

In the last equality, we used the fact that any constant can be added to y_n without affecting that sum, and in particular the factor $y_n - \mu_h$ can be replaced by $y_n - \mu_y$. This shows the relation to the estimated correlation function $\hat{R}_{hh}(x)$.

Note that the last term can be interpreted as the estimated variance of the terrain profile based on the available samples. Thus, the LS estimate is biased towards areas with flat terrain compared to the correlation approach.

However, the least squares approach is the optimal approach in the navigation model, since this is how the Kalman filter and other filters would handle the terrain altitude information (*i.e.* optimal fusion with respect to the measurement model). Note that the expression (5b) readily generalizes to an m -dimensional terrain.

B. Uncertain Geometry

We now consider the two other cases where the array geometry is not precisely known in an m -dimensional terrain, starting with the general case (iii). Since $h(x)$ is highly nonlinear and W is stochastic, Y is in general not Gaussian distributed. To find a closed-form expression for the distribution of Y a first-order Taylor expansion of every measurement y_n about w_n can be employed:

$$Y \approx \hat{Y} = H(x \oplus D) + H'_W(x \oplus D)W + E, \quad (6a)$$

$$H'_W(x \oplus D) = \nabla_W H(x \oplus (D + W))|_{W=0}, \quad (6b)$$

which leads to the approximate distribution

$$\hat{Y} \sim \mathcal{N}(\mu_{\hat{Y}}(x), \Sigma_{\hat{Y}}(x)), \quad (7a)$$

$$\mu_{\hat{Y}}(x) = H(x \oplus D), \quad (7b)$$

$$\Sigma_{\hat{Y}}(x) = H'_W(x \oplus D)\Sigma_W(H'_W(x \oplus D))^T + \Sigma_E, \quad (7c)$$

where ∇_x is the Jacobian of $h(x)$ with respect to x and is defined as

$$\nabla_x h(x) = \begin{bmatrix} \frac{\partial h_1}{\partial x_1} & \cdots & \frac{\partial h_1}{\partial x_m} \\ \vdots & \ddots & \vdots \\ \frac{\partial h_N}{\partial x_1} & \cdots & \frac{\partial h_N}{\partial x_m} \end{bmatrix}. \quad (8)$$

Since (6a) is based on a Taylor expansion about W , the quality of the approximate distribution will depend on how linear the terrain is around the measurement points and the size of the error terms W . Using the approximate distribution in (7a) we can apply the *weighted least squares* (WLS) method

$$\hat{x} = \arg \min_x V(x), \quad (9a)$$

where

$$V(x) = (Y - \mu_{\hat{Y}}(x))^T (\Sigma_{\hat{Y}}(x))^{-1} (Y - \mu_{\hat{Y}}(x)). \quad (9b)$$

The expression in (9) constitutes a possible estimator.

If we, again, instead consider a one-dimensional terrain and independent errors W (*i.e.*, case (ii)) then according to (6a) the measurement Y is approximately Gaussian distributed with mean and covariance given by

$$\mu_{\hat{Y}}(x) = H(x \oplus D), \quad (10a)$$

$$\Sigma_{\hat{Y}}(x) = \text{diag}(\hat{\sigma}_0^2, \dots, \hat{\sigma}_{N-1}^2), \quad (10b)$$

$$\hat{\sigma}_n^2(x) = (h'(x + d_n))^2 \sigma_{w,n}^2 + \sigma_e^2, \quad (10c)$$

where $\sigma_{w,n}^2$ is the one-dimensional equivalent of $\Sigma_{w,n}$. Using (10), the cost function of the WLS (9b) can be expressed as

$$V(x) = \sum_{n=0}^{N-1} \frac{(y_n - h(x + d_n))^2}{(h'(x + d_n))^2 \sigma_{w,n}^2 + \sigma_e^2}. \quad (11)$$

Evidently, measurements over steep terrain are weighted lower in the WLS when the relative geometry of the array is not fully known.

IV. CRAMÉR-RAO LOWER BOUND

The Cramér-Rao lower bound provides best case performance for estimation problems. Analyzing the CRLB gives insight into how difficult an estimation problem is, and can be used as guidance when selecting sensor configurations.

A. Background

The CRLB, P^{CRLB} , constitutes a lower bound of the variance of an unbiased estimator \hat{x} of x , given a measurement y with known likelihood function,

$$P^{\text{CRLB}}(\hat{x}(y)) \succeq (\mathcal{I}(x))^{-1}, \quad (12)$$

where $\mathcal{I}(x)$ is the Fisher information matrix (FIM), evaluated in the true parameter value. The CRLB is valid under mild regularity conditions, which, *e.g.*, guarantee that the likelihood $p(y|x)$ is twice differentiable. The FIM is given by

$$\mathcal{I}(x) = \mathbb{E}(\nabla_x^T \log p(y|x) \nabla_x \log p(y|x)). \quad (13)$$

In the case where $y \sim \mathcal{N}(\mu(x), \Sigma(x))$, which corresponds to (7a), the FIM can be expressed as

$$[\mathcal{I}(x)]_{i,j} = \frac{\partial \mu^T(x)}{\partial x_i} \Sigma^{-1}(x) \frac{\partial \mu(x)}{\partial x_j} + \frac{1}{2} \text{tr} \left(\Sigma^{-1}(x) \frac{\partial \Sigma(x)}{\partial x_i} \Sigma^{-1}(x) \frac{\partial \Sigma(x)}{\partial x_j} \right), \quad (14)$$

where $\text{tr}(A)$ denotes the trace of a square matrix A .

Once again, the expression (14) is hard to interpret, therefore we resort to the one-dimensional case to gain deeper understanding.

B. Known Geometry

We consider the measurement model corresponding to the case of a known geometry (i). Using the fact that $Y \sim \mathcal{N}(H(x \oplus D), \Sigma_E)$, the expression (14) simplifies to:

$$\mathcal{I}_n(x) = \frac{1}{\sigma_e^2} (h'(x + d_n))^2 \quad (15a)$$

$$P^{\text{CRLB}}(x) = \left(\sum_{n=0}^{N-1} \mathcal{I}_n(x) \right)^{-1}, \quad (15b)$$

where \mathcal{I}_n is the Fisher information associated with measurement n and where we have used the fact that the Fisher information is additive for independent measurements.

The expression in (15a) indicates that the slope of the terrain is important, which has previously been reported for single-sensor TAN [2]. We can also observe that the variance is expected to be linearly dependent on σ_e . This observation implies that the quality of the altitude sensor is important. Likewise, if we assume the derivatives $h'(x + d_n)$ to be in the same order of magnitude, *i.e.*

$$|h'(x)| \approx |h'(x + d_1)| \approx \dots \approx |h'(x + d_{N-1})|,$$

then we can approximately express the CRLB as

$$P^{\text{CRLB}}(x) \approx \frac{\sigma_e^2}{(h'(x + d_n))^2 N}, \quad n \in 0, 1, \dots, N-1,$$

meaning that the lower bound for the variance would involve the term $N^{-\frac{1}{2}}$. From this, we would expect that for a small number of sensors we will see a large improvement in terms of efficiency by adding more, but as the number of sensors increases the rate of improvement will become slower.

Note, the Fisher information and CRLB is a local property, which does not take into account the difficulty of differentiating between different peaks in the likelihood. Its usefulness in indicating expected performance of an estimator is therefore limited to the case when the true peak in the likelihood is discernible.

C. Uncertain Geometry

We now consider the case (ii), where the array geometry is uncertain but independent for each sensor, *i.e.* W is diagonal. Using the approximate Gaussian expression of $Y \sim \mathcal{N}(\mu(x), \Sigma(x))$ in (10), yields an approximate expression of the CRLB by using (14):

$$\mathcal{I}_n(x) = \frac{(h'(x + d_n))^2}{\sigma_e^2 + (h'(x + d_n))^2 \sigma_{w,n}^2} + 2 \left(\frac{\sigma_{w,n}^2 h'(x + d_n) h''(x + d_n)}{\sigma_e^2 + (h'(x + d_n))^2 \sigma_{w,n}^2} \right)^2, \quad (16a)$$

$$P^{\text{CRLB}}(x) = \left(\sum_{n=0}^{N-1} \mathcal{I}_n(x) \right)^{-1}, \quad (16b)$$

where we have again used the fact that the Fisher information is additive for independent measurements. Note that if $\sigma_{w,n}^2 = 0$ the expression in (15a) is recovered. We can expect that as long as $\sigma_{w,n}^2$ remains small, then the observations made in regard to (15a) will generally hold for (16a) as well.

Of note here is that the second term in (16a) indicates that a larger $h''(x + d_n)$ will lead to a more informative measurement. Perhaps more unintuitive, if $h''(x + d_n) \neq 0$, $\sigma_e^2 \neq 0$, then the second term seems to increase as $\sigma_{w,n}$ increases. The first term seemingly behaves more along what could be expected, becoming smaller as $\sigma_{w,n}$ increases in size. Admittedly, as was mentioned in Sec. III-B, the expression (6a) is based on a Taylor expansion about W , and is therefore only valid when $h''(x + d_n)$ and $\sigma_{w,n}$ are small.

Finally, for the general case (iii) we can also derive an approximate CRLB by plugging in the approximate mean and covariance $\hat{\mu}$, $\hat{\Sigma}$ from (6a) into (14). However, since there is no structure in Σ_W to exploit, we have not been able to find a simpler expression to analyze.

V. PRACTICAL MAP CONSIDERATIONS

To evaluate the terrain at intermediate points in the DEM it is necessary to either use interpolation or a more general representation. Traditionally in TAN, the map is assumed to be linear between the points in the DEM and so intermediate points are calculated using linear interpolation [2]. Since (16a) involves second order differentiation of the map, we would like a continuous representation that admits non-zero second order derivatives. One possibility which we have opted for in this paper is to consider cubic splines. We limit ourselves to

consider one-dimensional cubic splines here, but the concept can be generalized to surfaces [16].

Given a dataset consisting of M data points, constructing a one-dimensional piecewise cubic polynomial representation amounts to defining and solving a system of $4(M - 1)$ equations (there are $M - 1$ intervals with polynomials, each with 4 parameters). The typical approach is to define

- $2(M - 1)$ equations specifying end point values as data point values,
- $M - 2$ equations specifying continuous first order derivatives at the “inner data points”,
- $M - 2$ equations specifying continuous second order derivatives at the “inner data points”,

This still leaves two equations to specify. In the *natural approach*, the second derivative of the end points is set to zero. This is very simple, but may not always be very accurate [17]. Another simple approach which typically produces more accurate results is the *not-a-knot* end conditions, where the third order derivative of the end points is set to be equal to their closest data point [17].

Using this approach, a one-dimensional piecewise cubic polynomial terrain can be constructed by solving the above systems of equations using the points in the DEM. By doing so, the first and second order derivatives of the terrain have an analytical expression.

VI. NUMERICAL RESULTS

Numerical simulations will now be used to illustrate and verify the properties derived in Sec. III and Sec. IV.

A. Simulation Setup

As a part of the simulation study, we consider a variant of case (ii) presented in the text. Specifically, we consider N sensors measuring the altitude over a one-dimensional terrain $h(x)$ with an error in the relative positions modelled as independent and identically distributed $W \sim \mathcal{N}(0, \sigma_w^2 I)$. For simplicity, we consider the mean relative positions between the platforms to be an equidistant array, *i.e.*, $d_n = nd$, $n = 0, \dots, N - 1$. This will not have an effect on the results beyond potential minor alias effects.

The terrain used in the simulation study, which is shown in Fig. 2, consists of samples from a real DEM of Sweden with resolution $1 \times 1 \text{ m}^2$. The piece of terrain is 2000 m long, and we represent it by constructing a cubic spline representation with the not-a-knot conditions as described in Sec. V. Since the samples in the DEM are so dense, the resulting spline will be representative of the true terrain. We consider the problem of locating the first sensor in the array somewhere over the terrain using the WLS estimator defined in (9). We assume no prior knowledge about the position of the first sensor, corresponding to a uniform prior.

We are interested in evaluating both the robustness of the estimator and its efficiency. In this context, we refer to the robustness of the estimator as its ability to discern the true peak in the likelihood. A simple measure of this is the fraction of successful estimates (discerned peaks) by the estimator over




Fig. 2: Terrain curve with resolution 1 m from the south of Sweden taken from a terrain database. The original terrain database consists of a grid with resolution $1 \times 1 \text{ m}^2$. Markh jdmmodell Nedladdning, grid 1+.  Lantm teriet

a number of evaluations. We consider a simple condition to check if a trial was successful by defining a threshold value δ and consider errors outside this value failures:

$$\begin{cases} \|\hat{x} - x\| \leq \delta, & \text{success,} \\ \|\hat{x} - x\| > \delta, & \text{failure,} \end{cases} \quad (17)$$

where $\|\hat{x} - x\|$ is the estimation error of an estimate \hat{x} given the true value x . To evaluate the efficiency, we consider how closely the *root mean square error* (RMSE) of the estimator matches the CRLB. To separate the effects of robustness and efficiency, we only include successful trials in the RMSE, and which will correspond to a quantity that matches closer to the CRLB which is conditioned on the true parameter value.

For evaluation purposes, the terrain is divided into a grid of size 1 m (*i.e.*, 2000 grid points). At each grid point, the CRLB is calculated using (16a). The resulting CRLB values are then averaged over the grid points to yield a single value. The estimator is evaluated by, at each grid point, calculating snapshot estimates using Monte Carlo trials. In the Monte Carlo trials the errors W and E are the random quantities. Each trial is labeled a success or failure using the threshold value, $\delta = 50 \text{ m}$ as in (17), and which yields the fraction of successful trials. The results from the successful Monte Carlo trials are used to calculate the RMSE for a snapshot estimate over the terrain using the expression

$$\text{RMSE} = \sqrt{\frac{1}{M} \sum_{m=1}^M \frac{1}{N} \sum_{k=1}^K \|x_k - \hat{x}_{k,m}\|^2}, \quad (18)$$

where x_k is grid point k , $\hat{x}_{k,m}$ is the estimated position at grid point k during the successful trial m , $K = 2000$ is the number of grid points in the map and M is the number of successful MC trials. The average CRLB over the terrain P^{CRLB} compared to the RMSE e^{RMSE} and the fraction of successful trials together constitute the performance of the estimator for one set of parameter values.




Fig. 3: Results for tests when varying the standard deviation of the error in the relative position between the sensors σ_w . Here, P^{CRLB} is the full CRLB expression (16a), and P_1^{CRLB} is the same CRLB expression but where the second term has been omitted. The fraction of successful snapshots varies from 99% when $\sigma_w = 0$ to 98% when $\sigma_w = 9$.

B. Simulation Results

We examine the properties of the estimator by considering a nominal set of parameters

$$\sigma_e^2 = 1, \sigma_w^2 = 25, N = 10 \text{ and } d = 20,$$

and then vary each parameter. In doing so, the effect of each parameter on the estimation performance can be examined separately.

In Fig. 3 σ_w is varied. Here, both the CRLB expression containing only the first and both terms in (16a) are plotted, denoted P_1^{CRLB} and P^{CRLB} respectively. We can see that the inclusion of the second term leads to a lower theoretical variance. Clearly, the expression where the second term has been omitted is more representative of the RMSE and is more intuitive, as it increases when σ_w increases. Due to this result, P_1^{CRLB} is used instead of P^{CRLB} throughout the rest of the paper. The fraction of successful trials as presented in Fig. 3 seems to only be affected slightly as the SNR decreases.

In Fig. 4 the mean distance d between the platforms is varied. We can see that increasing d seems to have a small positive effect on the RMSE and number of successful trials at first, which then tapers out. A plausible interpretation becomes apparent if we consider points close to each other in the terrain to be correlated. Increasing the distance will then have a positive impact until the points become “uncorrelated”.

The number of platforms N is varied in Fig. 5. Here we see a sharp decrease in the error and fraction of unsuccessful trials as additional sensors are added for a low number of sensors. Adding more sensors to an already large amount of them seemingly still has a modest positive effect. The trend seems to be in line with the expected results in Sec. IV-B.




Fig. 4: Results for tests when varying the mean distance between the sensors d . We can observe a moderate decrease in the error and number of unsuccessful snapshots, which tapers out at around $d = 40$.

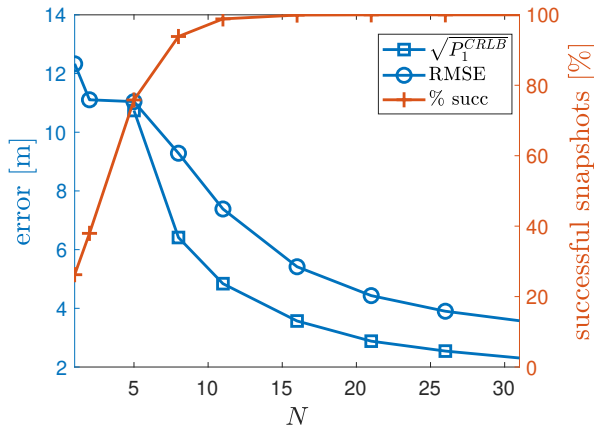


Fig. 5: Results for tests when varying the number of sensors N . Here the values of $\sqrt{P_1^{CRLB}}$ are omitted for $N < 5$. This is due to the fact that using a low number of sensors easily leads to at least one grid point in the map where the sensor array measures featureless points, *i.e.* where $h(x) \approx 0$. This leads to a very small value of the Fisher information, and in turn a very large CRLB value. Since $\sqrt{P_1^{CRLB}}$ is the average CRLB over the grid, even one such point leads to an unrepresentative average.

Finally, in Fig. 6 the variance of the error in the sensors σ_e is varied. The error and fraction of unsuccessful trials seem to grow rapidly as σ_e increases. This result is expected, since if the terrain is not observable in the noise then increasing the number of measurements will only have a small impact. It is also in line with the discussion in Sec. IV-B, with $\sqrt{P_1^{CRLB}}$ seemingly increasing linearly with σ_e . On the other hand, the RMSE does not follow the same trend as $\sigma_e > 2$. This can be understood as the variance of the estimator becoming very large, so that even “legitimate” outcomes of the trials are




Fig. 6: Results for tests when varying the standard deviation of the altimeters σ_e . As σ_e increases, the CRLB, RMSE and % of unsuccessful snapshots rapidly increase.

labeled as failures, resulting in the seemingly stable value of the RMSE above $\sigma_e \geq 3$.

VII. CONCLUSION AND FUTURE WORK

This paper introduced multi-agent *terrain-aided navigation* (TAN) as a method for navigation in GNSS-denied environments. Contrary to the single-agent TAN, which is a fairly well studied topic, multi-agent TAN is less studied. The multi-agent setup TAN problem has been properly formulated, and its properties were studied in the context of altitude-based TAN. The *Cramér-Rao lower bound* (CRLB) was derived, and the effect of varying the different parameters was analyzed. This analysis was then complemented with a simulation study to quantify the results. One interesting observation that was made is that the solution is relatively insensitive to errors in the relative positions of the agents, which indicates that relatively low accuracy inertial navigation systems and distance sensors should suffice when determining the agent configuration. The CRLB and the results indicate that there is a lot of performance to be gained by increasing the number of agents beyond one, but that the effect is less significant asymptotically. However, it was also observed that to increase the number of agents is not a solution to low quality altitude sensors. That is, adding more agents is not a remedy for poor altitude sensors. It was also shown that to be efficient, the agents must be separated in space to utilize the variations in the map, but that once enough separation is obtained, separating the agents more will not improve the localization performance further. Combined, all this provides guidelines for how to design multi-agent TAN systems, and at the same time highlights aspects of the concept that need to be explored further.

For future work we would like to consider the connection between the terrain and the performance closer. By considering the terrain as a realization from a stochastic process, it should be possible to connect the findings in this paper to the properties in the terrain, such as the correlation.

ACKNOWLEDGMENT

This work was performed within the Competence Center SEDDIT (Sensor Informatics and Decision making for the Digital Transformation), supported by Sweden's Innovation Agency within the research and innovation program Advanced digitalization.

REFERENCES

- [1] Daniel Medina, Christoph Lass, Emilio Pérez Marcos, Ralf Ziebold, Pau Closas, and Jesús García. On GNSS Jamming Threat from the Maritime Navigation Perspective. In *2019 22th International Conference on Information Fusion (FUSION)*, pages 1–7, 2019.
- [2] Niclas Bergman, Lennart Ljung, and Fredrik Gustafsson. Terrain navigation using Bayesian statistics. *IEEE Control Syst. Mag.*, 19(3):33–40, June 1999.
- [3] Aaron Canciani and John Raquet. Airborne Magnetic Anomaly Navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1):67–80, 2017.
- [4] Peter F. Swaszek, Richard J. Hartnett, and Kelly C. Seals. Celestial Navigation - Correcting the Folklore. In *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 949–959, 2020.
- [5] Fredrik Gustafsson, Fredrik Gunnarson, Niclas Bergman, Urban Forssell, Jonas Jansson, Rickard Karlsson, and Per-Johan Nordlund. Particle filters for positioning, navigation, and tracking. *IEEE Trans. Signal Process.*, 50(2):425–437, February 2002.
- [6] Rickard Karlsson and Fredrik Gustafsson. Bayesian surface and underwater navigation. *IEEE Trans. Signal Process.*, 54(11):4204–4213, November 2006.
- [7] Per-Johan Nordlund and Fredrik Gustafsson. Marginalized particle filter for accurate and reliable terrain-aided navigation. *IEEE Trans. Aerosp. Electron. Syst.*, 45(4):1385–1399, October 2009.
- [8] I. Nygren. A method for terrain navigation of an AUV. In *MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings (IEEE Cat. No.01CH37295)*, volume 3, pages 1766–1774 vol.3, 2001.
- [9] André Przewodowski, Fernando Santos Osório, and Valdir Grassi Junior. Global Localization using OpenStreetMap and Elevation Offsets. *Journal of the Brazilian Computer Society*, 30(1):264–273, Sep. 2024.
- [10] Nicolas Souli, Rafael Makrigiorgis, Panayiotis Kolios, and Georgios Ellinas. Cooperative Relative Positioning using Signals of Opportunity and Inertial and Visual Modalities. In *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, pages 1–7, 2021.
- [11] Nikos Piperigkos, Aris S. Lalos, Kostas Berberidis, and Christos Anagnostopoulos. Cooperative Multi-Modal Localization in Connected and Autonomous Vehicles. In *2020 IEEE 3rd Connected and Automated Vehicles Symposium (CAVS)*, pages 1–5, 2020.
- [12] Pengxiang Zhu, Yulin Yang, Wei Ren, and Guoquan Huang. Cooperative Visual-Inertial Odometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13135–13141, 2021.
- [13] André R. Braga, Marcelo G. S. Bruno, Emre Özkan, Carsten Fritsche, and Fredrik Gustafsson. Cooperative Terrain Based Navigation and coverage identification using consensus. In *2015 18th International Conference on Information Fusion (Fusion)*, pages 1190–1197, 2015.
- [14] Hallysson Oliveira, Stiven Schwanz Dias, and Marcelo Gomes da Silva Bruno. Cooperative Terrain Navigation Using Hybrid GMM/SMC Message Passing on Factor Graphs. *IEEE Transactions on Aerospace and Electronic Systems*, 56(5):3958–3970, 2020.
- [15] Hallysson Oliveira, Stiven S. Dias, and Marcelo G. S. Bruno. GNSS-Denied Joint Cooperative Terrain Navigation and Target Tracking Using Factor Graph Geometric Average Fusion. *IEEE Transactions on Aerospace and Electronic Systems*, 60(1):991–1005, 2024.
- [16] David B. Kidner. Higher-order interpolation of regular grid digital elevation models. *International Journal of Remote Sensing*, 24(14):2981–2987, 2003.
- [17] Carl d. Boor. *A Practical Guide to Splines*. Springer Verlag, New York, 1978.